

Coup Agents: Agents that play a deception game

(extended abstract of the MSc dissertation)

Filipe Fortes
Departamento de Engenharia Informática
Instituto Superior Técnico

Abstract— This work investigates how to build an agent that, by being implemented in a robotic entity, is capable of playing a deceptive social game called COUP at the same level of humans. To be able to do so, we first started by defining the overall problem followed by how our solution should be. We then produced an algorithm to make the decisions for the COUP game, which is based on minimizing counterfactual regret. A system based on the before mentioned solution was implemented, so that we could test our hypothesis. This system includes our agent's architecture, the game interface and EMYS, which is the robot we used to physically represent our agent. To prove our hypothesis we devised an experiment with four different conditions, a group and individual condition, and a lie and truth condition. In every condition, our agent was capable of playing at the same level of humans, with the exception of the individual lie condition, where he was slightly better at playing the Coup game than its human opponents. This proves that it is indeed possible for an artificial agent to play a deception game at the same level of humans.

I. INTRODUCTION

Deception has been used by various types of life forms to achieve different objectives[3], for example, chameleons use deception as a defense mechanism. While usually carrying a negative connotation, many studies show that it appears everyday in our lives[8], ranging from little lies, such as telling someone that he/she can have our dessert because we are full, when we are not, so that the person would not feel guilt for eating it, to bigger, more dangerous lies, that is a lie about infidelity.

Since deception is so commonly used among humans, it should be taken into consideration when building a robotic agent capable of social interaction. But that is not the sole reason that it should be considered, if we go into the most general definition of the word, “a false communication that tends to benefit the communicator”[3], we can see that by adding deception into the architecture of such robotic agent, we open a whole new world of possible outcomes for such entity. An example of this would be a situation where a robotic agent tells a criminal that he is looking for the keys of the safe, while in reality is waiting for the police to reach his position.

However, it is not always easy to find such situations where one can lie multiple times without having social repercussions. A game constitutes the perfect option for studies about deception, because not only it allows a rich environment for deception, but also provides situations where one can lie without being affected by it on his/her social life.

Coup is a turn-based social game, that can be played by two to six people, where one uses deception to maintain his/her influence (face down cards) while making other players lose theirs. The deception used in this game is mainly towards what are the cards that each player has face down, making it a great game to build a deceptive robotic agent upon, since communication options are limited but still allows for a great deal of deception.

Using the Coup game, we address the general problem of creating deceptive behaviors in a robot or an agent. More concretely, we address the following question: “Is it possible to create an agent that, by implementing it on a social robot, is capable to play the Coup game at the same level of humans?”. Which in turn, we predict that it is indeed possible to do so.

To answer the before mentioned question and to actually know if our prediction is indeed true, we have built an agent architecture that takes into account the possibility of deception and uses a decision making algorithm to produce actions for the Coup game. We have also built a complete system that our agent can use to both play the Coup game against human players and interact with them through the use of a social robot.

II. BACKGROUND

A. Deception

Deception can be a mean to an end but it usually carries a negative connotation among humans. Nevertheless, the animal kingdom is filled with all kinds of deception, ranging from mimicry and camouflage, to the more commonly seen feigning death. All these mechanisms, actively or passively used by animals, have been developed during the evolution of the species and is what allows a big part of the animal kingdom to survive against their predators.

This shows that deception, at least when used in the animal kingdom, represents an evolutionary advantage for the deceiver. Further, some researchers even point out that deception is a strong potential indicator of theory of mind[6] and social intelligence[12].

So what is deception? Along this work we will be using the definition for deception provided by Bond and Robinson, and they describe it as “a false communication that tends to benefit the communicator”[3].

When applied directly to the human being, deception can have various types. These types include: (1) intentional and accidental; (2) active and passive deception; (3) implicit and

explicit deception; and finally (4) strategic deception, for instance, when a poker player checks instead of raising the bet, so that the other players will think that his hand is not as strong as it really is.

Ways of detecting deception pass by detecting the signals usually sent by the deceiver when deceiving. These include making facial expressions that differ from what the person is trying to say that he/she is feeling, as mentioned before; the deceiver's body language and even the volume of his/her voice and the pauses throughout the speech.

B. Board Games

A board game, in its most general definition, is a game that involves the placing or moving of some kind of objects on a pre-marked surface called "board", according to a given set of rules.

The real magic to these kinds of games is that the players, not only are having fun when playing the game, but are also learning how to deal with situations that are not so common in their normal lives and gaining skills that would be much harder if learnt in another way.

Deception board games encourage the players to use deception to achieve their aims and always have an element of hidden information in them. They are usually linked to party games as they encourage social interaction, and most of the deception types require some sort of interaction between people. The game Werewolf and its variant Mafia are one of the most known deception games. It is a game where there are two parties, the werewolves and the townspeople, with differing objectives: the werewolves want to kill the townspeople and the townspeople the werewolves. The catch in this game is that the townspeople don't know who the werewolves are, so the werewolves have the objective of blending in with them, using deception to do so.

Another game that resembles the Werewolf game is The Resistance game, where the werewolves are the spies and the townspeople are the resistance. The mechanics are slightly different so that no player elimination happens during the game. Coup is a game from The Resistance's family, that provides a completely different gameplay, while still allowing a great deal of deception.

C. COUP

Coup is a social board game that revolves around secret identities, deduction and deception. The game consists in making other players lose their influence (cards on the board) while not losing yours, or at least be the only player with influence at the end of the game. Each influence represents a character that has some abilities and only two influence is given to each player at the starting of the game. When a player loses all its influence, he is exiled and loses the game.

The game is played in turns and in each turn a player chooses an action from the list of actions, after that, any other player has the ability to challenge or counteract that action. If the action is challenged, the player that tried to perform the action must prove that he has the card that can execute such action, if he proves it by revealing the card, the

challenger will lose one influence, otherwise it will be the challenged player that loses one influence. When a player performs a counteraction, that too can be challenged and if it succeeds the original action will be considered void and that player's turn will be expended. It is also important to mention that only three cards of each character exist.

The actions revolve around getting coins to be able to expend them on actions that make other players lose influence, such as the "Assassinate" and "Coup" actions.

Coup provides one of the best blends of deception and board games, while still providing a simple rule set and well defined actions, making it a great game to base our work on.

III. RELATED WORK

This section has the objective to show some of the work that has already been done in the area of agents with deception. While they do not focus directly on using such deception for social board games, which is our objective, they are the closest to it in today's state-of-the-art.

A. GOLEM

GOLEM is a system, developed by Castelfranchi, Falcone and de Rosis[5], based on a multi-agent blocks world with the objective of studying both the interactions and attitudes between two agents with different social attitudes and personalities, when delegating and adopting tasks from each other.

GOLEM's multi-agent world only contains two agents, where each of those agents has a different goal and tries to achieve it. Personalities of the agents is what makes GOLEM such a rich environment, as they are much more diverse. Both agents have personalities for when delegating tasks, from "Lazy Agent", to "Never-Delegating" ; and when adopting tasks, from "Hyper-Cooperative" , to "Non-Helper". Each agent is also limited in terms of abilities, or in other words, each agent may only be capable of performing a limited set of actions on the domain state.

GOLEM acts as a game where initially both agents will introduce themselves by stating their personalities and abilities, which they can lie about, or simply give partially incorrect or imprecise descriptions. The running of this simulation is then played in turns, where each turn an agent can perform some action on the domain or not, and perform a "communicative act" according to a defined protocol.

Agents in GOLEM are able to deceive about their capabilities, their personality and even their goals and plans, and can do so in the situation that if the other agent knows about their true properties, it will not adopt the request made by him.

While GOLEM allows more kinds of deception than any of the other related works, it still does not allow deception through speech acts, which may have value if implemented in our work. Nevertheless, GOLEM still provides a relatively better mental model modulation that may be similar to what we need, and it also provides an overall good basis for our modulation of deception.

B. Deception Planner

Deception Planner, developed by David Christian for his master thesis[7], is an implementation of a model of strategic deception, i.e. it attempts to deceive in order to achieve or enable some final goal.

The problem that the Deception Planner has to deal with is to find which statements, that may include lies, should an agent present to another agent in a way that the actions of the second agent will achieve the ulterior goals of the first agent. To do so, it first needs some input of the problem, that is, a set of ulterior goals, the current world state and the model of target agent, which include observation rules. In the end, it will provide, if possible, a set of facts and a set with negation of facts.

The deception planner is a modified version of the LPG (Local Search for Planning Graphs) planner, for the reason that it is a local search planner, the heuristic is relatively informed and is able to do plan repair.

To be able to find a plan that fulfills the needed conditions, the deception planner uses a complex heuristic that includes the generic costs used in the generic LPG planner, CostToFixMutexes, the cost of reasserting a condition due to a mutex (mutual exclusive actions), and the CostToFixOpen, an estimation of the cost of achieving all open preconditions of the new plan, while also adding CostToAchieveUltGoals, which is an estimation of how many steps are needed from the current plan to a plan that achieves all ulterior goals, and CostToFixLies, which is an estimate of how many steps have to be added so that the lies are kept from being observed.

The last part of the deception planner is the negation of competing plans. A competing plan is a plan that is better for the target agent and does not achieve the ulterior goals. To negate such a plan, the deception planner tries to find lies about beliefs that the target agent may have, so that those competing plans are no longer viable.

The Deception Planner is a very nice tool for deception and is a potential addition for our agent, so that it can produce deception through speech acts.

C. Study about Robots Deceiving Humans by Terada and Ito

Terada and Ito have developed an experiment to prove if robots can indeed deceive humans. While they do not provide a tool of some sort towards deception, like the majority of related works, their work has value when we take into consideration the true meaning of deception and try to apply it to a robot[16].

To prove that robots can deceive humans, the authors focus on deception as a cue for deception attribution and base their experiment on proving that a human can treat a robot as an intentional entity.

The experiment consists of two phases: (1) the robot is facing the wall and says “Daruma-san go Koranda”, while the person is walking towards the robot, and then turns around, and (2) tries to detect if the person is moving or not and turns to face the wall again. These phases are repeated until either the robot detects the person moving or the person is able to touch a button on the robot’s head surface.

To be able to reach some conclusions, they used two different experimental conditions, (1) the deception condition, where the robot adopted a “slow normal behavior” which consisted on slow chanting of the syllables and a big turn around time, and then, when the person could reach the robot in the next turn, it would adopt a “fast deception behavior”, accelerating the chanting of words and turn around; and (2) the control condition, where it always produces the “slow normal behavior”. The robot in the deception condition always signalizes that he caught the person moving, even if that is not true.

They concluded, through the analysis of questionnaires given to the participants at the end of the experiment, that the deception condition group felt more outwitted than the control condition group, which in turn can be concluded that the robot was perceived as an intentional entity.

This work and its approach is quite simple and makes some bold assumptions. Nevertheless, it provides a strong indication that robots can indeed deceive humans, which is a good assumption to have in our work.

D. Mindreading Agents

João Dias et al. developed a model for a mindreading agent that supports N levels of Theory of Mind and is capable of carrying out deceptive behaviors[9].

Their agent’s Theory of Mind is based on the Mindreading model of Baron-Cohen[1] and follows the ST (Simulation-Theory) of Meyer et al.[11], which claims that one should represent others the same way they would simulate themselves in the same situation.

They built an experiment to test if agents with two levels of ToM were more effective than agents with only one level of ToM, when playing a deceptive game called Werewolf. The results showed that the 2-level ToM agent version won more games than the 1-level ToM version, proving that having it is advantageous for an agent to possess two levels of ToM when playing a deceptive game as opposed to only one level.

The work done by João Dias et al. is strongly connected to ours and provides us with a solid prove, given this work’s subject, that having two levels of theory of mind gives an advantage to an agent capable of deceiving, over only one level of theory of mind.

IV. PROBLEM AND SOLUTION DEFINITION

In this section we start by giving a detailed description of the problems that are inherent to the problem that this work solves, and provide an overview of the most important aspects to take into consideration for the solution that is presented after. We then proceed to describe the various components of the solution devised to solve the problem.

A. Problem

The first part of the problem that we address is the need for our agent to be able to controllable how it lies.

The agent needs to be aware of the drawbacks of lying, since without it, it would perform a variety of actions that would make sense because of the value they would have in

some particular moment and not because those are the actions that he is able to do. By doing so, the agent would probably be constantly lying and by doing that, it would keep losing indefinitely, as it is much easier to catch when the agent is lying.

To be able to simulate a human player in terms of behaviour and to give the feeling that the agent is not completely virtual (give feeling of familiarity to the humans towards the agent), we need some kind of physical entity to represent our virtual agent. Having such entity would allow the virtual agent to interact with the physical world, making it capable of creating a more solid relationship with its fellow human opponents.

We also need to consider that our agent has to be capable of using the physical entity capabilities in order to take advantage of said interactions.

In terms of game interfaces, to be able to test an agent that is capable of deception for its own advantage while playing the board game COUP, we have to have an environment where both the virtual agent and the physical human can play the game simultaneously.

To make this work the best possible, an equilibrium between a virtual game and the original physical board game has to exist. Something that feels like the traditional physical game and that provides some type of interface for the virtual agent, so that the agent can perceive what is happening in the game and at the same time, perform actions on it.

To effectively demonstrate some similarities to the way that humans behave, the virtual agent and the physical entity that represents the agent have to be linked so that when the virtual agents decides some action towards the game, the physical entity will represent how the agent would behave while doing that action. This implies that the agent is capable of sending commands to the physical entity, or that the physical entity is capable of recognizing the actions made by the agent and acted in conformity to that.

Being an artificial entity, our agent will not have the same capabilities of a human in terms of lie detection. Humans subconsciously detect out of order patterns on other humans. These can be changes in facial expressions, changes in vocal intensity, stuttering, changes in body expressions. A virtual agent is not capable of detecting these. It would require some good cameras capable of detecting those changes and software that could determine if those expressions were coming out of the ordinary. This is not possible with the equipment we have available to us.

Finally and mainly, the most significant problem that we have to solve is to make our agent capable of playing at the same level of humans. To do so would imply that, not only is our agent capable of detecting lies at the same level of humans, but also make decisions towards what actions it should take, including actions that it should not be able to do considering the character cards it has, or in other words, perform actions that require lying. This problem is not as linear as we are putting it, since by having a very strong decision making algorithm, our agent can have a weaker lie detection system and still balance things out in a way that it

is has capable of winning the game as humans are.

B. Solution

To solve the problem mentioned before, we have devised a system that is divided into three big parts:

1) Digital Tabletop: To meet all the requirements imposed to the game, regarding the different interfaces that it should have, we came to the conclusion that a digital tabletop would be our best bet, since it is capable of running the game as a software, while displaying it for the humans through a touchable digital display, much like the one used for smartphones and tablets, but the size of a table, and displaying it through a virtual interface for the virtual agent to receive messages of what is happening to the game state and send commands to execute actions unto the game. By having a physical touchable display and being capable of detecting and distinguishing objects that are on top of it, the digital tabletop provides a seemingly traditional board game experience for the human players.

Since most of the digital tabletops commercially available run the same operating systems as the ones that personal computers run, the game can be implemented in Unity. Which also provides the capacity to use both interfaces that we need, the one for the virtual agent and the one for the human players, putting it as one of our main building blocks for our work.

2) Robot: The solution we came up for the problem of the physical entity is to use a robot with all those capabilities. The robot has the form of a human head, which is capable of expressing a complete range of human emotional and intensity of such emotions and mobile in the same way a human head his, so that it can change the directionality of its focus, making the humans recognize to whom the agent is looking at or directing its interaction to. To meet the requirement of recognizing the target of interaction, the robot has a camera attached to itself, so that it can map the presence of each human in terms of directionality of view and then focus a specific human based on its objective of interaction. Another of the problems mentioned before was the need to make the agent able to express itself through voice output, for that, the robot must have a speaker also attached to it, so that directionality of interaction still exists. Finally and as a way to make the virtual agent capable of detecting the human voice, the robot needs to be connected to microphones, which are best used if attached to each human player, as the voice recognition and distinction will be much more easier to do and complex sound analysis software is not required.

With all this functions available to the virtual agent to express itself in the real world and interact with the humans, it is now provided with the much needed elements to achieve a deceptive behaviour and successfully deceive its human opponents.

3) **Agent's Architecture:** The last part is the most important one for our work, since it is the very core of our virtual agent. The agent's architecture part is responsible for all of the logic and processes that happen with the our agent, which include the agent-game communication and the agent-robot communication, this means that the architecture takes into consideration both the robot and the digital table interfaces to link them to the virtual agent.

An overview of the components that should be incorporated in our agent architecture is shown next:

- **Perception Receiver** It's the component responsible for receiving perceptions, both from the game and from the robot, for example when an interaction as finished.
- **Memory** Contains both the memory of the agent, which holds all the perceptions received, and its theory of mind, which takes into consideration the agent's memory and tries to derive what the opponents are thinking, including the probability of having each card and the probability of making a certain action.
- **Theory of Mind** Takes into consideration the agent's memory and tries to derive what the opponents are thinking, including the probability of having each card and the probability of making a certain action.
- **Decision Making Algorithm** Algorithm that uses the theory of mind component and a modification of the regret minimization algorithm to produce the next action for the agent to play
- **Action/Interaction Producer** This component is responsible for the decision of how the robot used to represent the agent's physical presence should act, it takes into consideration the action that the agent is going to make. It is also the connection that outputs messages to both the robot and the digital table so that actions can be played in both of them.

V. THE DECISION MAKING ALGORITHM

In this section we describe how the decision making algorithm works in detail. We first start to describe the standard algorithm on which the decision making algorithm is based on, which is the regret minimization algorithm in games with incomplete information. We then describe the necessary modifications to this algorithm in order to take into consideration some aspects that are game-dependent.

A. Regret Minimization

The decision making algorithm used in this work is a modification of the regret minimization algorithm that minimizes regret through the minimization of counterfactual regret [18]. The regret minimization algorithm not only supports games with incomplete information, which is the case of COUP, but also works fairly well in extensive games.

This algorithm was chosen as the starting point for its way of dealing with extensive games with incomplete information, as mentioned before, and for its amazing success in Poker, which is a similar game to COUP, in the way that a player, to be successful, will most likely have to lie.

To define the concept of regret, we need to consider playing an extensive game on a repeated way. Letting σ_i^t be the strategy used by player i on round t . We can calculate the average overall regret at time T of a player i with:

$$R_i^T = \frac{1}{T} \max_{\sigma_i^* \in \Sigma_i} \sum_{t=1}^T (u_i(\sigma_i^*, \sigma_{-i}^t - u_i(\sigma_i^t))) \quad (1)$$

The fundamental idea of the regret minimization algorithm mentioned before, is to decompose the overall regret into individual regret terms that can be added, which will then make them able to be minimized independently.

To do so, Zinkevich et al. first came to the most important key result from their approach. That $R_i^T \leq \sum_{I \in \mathcal{I}_i} R_{i,imm}^{T,+}(I)$. And so we actually know that by minimizing immediate counterfactual regret, we can minimize the overall regret.

To minimize the regret in an independent way for each information set, Blackwells algorithm for approachability can be used:

$$R_i^T(I, a) = \frac{1}{T} \sum_{t=1}^T \pi_{-i}^{\sigma^t}(I) (u_i(\sigma^t|_{I \rightarrow a}, I) - u_i(\sigma^t, I)) \quad (2)$$

Define $R_i^{T+1}(I)(a) = \max(R_i^T(I, a), 0)$, then the strategy for time $T + 1$ is:

$$\sigma_i^{T+1}(I)(a) = \begin{cases} \frac{R_i^{t,+}(I, a)}{\sum_{a \in A(I)} R_i^{t,+}(I, a)} & \text{if } \sum_{a \in A(I)} R_i^{t,+}(I, a) > 0 \\ \frac{1}{|A(I)|} & \text{otherwise.} \end{cases} \quad (3)$$

In other words, actions are selected in proportion to the amount of positive counterfactual regret for not playing that action. If there is no action that produces a positive counterfactual regret, then the action is selected randomly.

B. The Decision Making Algorithm

Coup is a game where the information available is incomplete, which leads to the need of making actions and decisions without having all the relevant information to get the best result out of the actions and decisions available at that moment. This will produce regret in a player that makes an action and then, afterwards, gets the remaining information that was missing and realises that if he would have chosen another action, he would have gotten a better result. By minimizing the counterfactual regret and, in turn, the overall regret of the agent, we make it possible for the agent to have a decision making algorithm that thrives in an environment where information is not perfect.

When calculating the utility of an information set given a strategy using $u_i(\sigma, I) = \frac{\sum_{h \in I, h' \in Z} \pi_{-i}^{\sigma}(h) \pi^{\sigma}(h, h') u_i(h')}{\pi_{-i}^{\sigma}(I)}$, we changed it so we never take into consideration specific histories, but just continue to deal with information sets, as the abstractions are still needed at this level.

We then came up with the following counterfactual utility $u_i(\sigma, I)$ function:

$$u_i(\sigma, I) = \sum_{I' \in (I, \sigma), I'' \in Z} \pi_{-i}^\sigma(I, I') \pi^\sigma(I', I'') u_i(I'') \quad (4)$$

Where $\pi_{-i}^\sigma(I, I')$ is the probability of state I' being the outcome of the current state I given that the player plays accordingly to σ and $\pi^\sigma(I', I'')$ is what we have called the *potentialToWin* which returns the estimated probability to win of the player. Finally, we only take into consideration the outcomes that make the player victorious so, $u_i(I'') = 1$ if the player is victorious and $u_i(I'') = -1$ if not. This allows us to remove a great amount of possibilities that need to be calculated.

In a way to simplify the function, but still reproducing the same expected results, we went a little further and modified $\pi^\sigma(I', I'') u_i(I'')$ into *potentialToWin*(I', σ), which removes all the need to have an utility function for each terminal state, since the ones where the player loses would account to -1 utility and the ones where he wins account to 1 utility. This *potentialToWin*(I', σ) function calculates the potential of the player to win the current game and returns a value between -1, game is already lost, to 1, game is already won.

The final counterfactual utility $u_i(\sigma, I)$ function used by our decision making algorithm is:

$$u_i(\sigma, I) = \sum_{I' \in (I, \sigma), I'' \in Z} \pi_{-i}^\sigma(I, I') \text{potentialToWin}(I', \sigma) \quad (5)$$

The modifications made to the algorithm only affect the way that the probability to a terminal state is calculated and with that the utility of the final state. In other words, the change only modifies how $\pi^\sigma(h, h') u_i(h')$ is calculated, but since *potentialToWin*(I', σ) still provides the same results expected from the generic algorithm, values between -1, if the game is lost, and 1, if the game is won, we can conclude that the our modified version of the algorithm will still provide the same properties as the original algorithm.

The strategies that correspond to the other players, we obtain them from the Theory of Mind component of our agent's architecture, so that all the Decision Making Algorithm can functional as intended.

VI. IMPLEMENTATION

In this section we will present how we actually implemented the more general solution that we defined before. We will also mention how everything works from an individual standpoint and how everything works together to produce exactly what we want.

A. Overall System

This system was built with the purpose of testing our hypothesis and for that, we had to take into account a great number of things. Starting with the core of the system, which is Thalamus, this components objective is to receive messages from the different components and send them to the components that are expecting to receive it. To do so,

Thalamus is composed by a scheduler integrated with a MOM (Message-oriented middleware), which allows for it to have asynchronous and abstract sides of communication, while still supporting synchronously distributed behaviours that will run in a BML-like manner. Since the Thalamus scheduler is more abstract than BML, it will allow the use of synchronized actions and events that are originated from BML-based behaviour [17]. This allows for the sending and receiving of events which is a good way to send information essentially between the game and our agent.

Our agent architecture, which is basically the core of our agent, will decide which actions and interactions should our agent perform, depending on the various situations that it will perceive. It is directly linked to a thalamus bridge so that it can convert the messages from Thalamus and filter them depending on which ones it wants to receive.

Similar to our agent architecture, the Coup Game component also has a thalamus bridge for the same exact reasons as the architecture. The coup game will send the perceptions that our agent will receive, as well as receive actions from our agent to apply into the game, and consequently show the human players how our agent as acted in the game.

The EMYS component represented in our overall system, not only includes the physical robot that will interact as the agent with the human players, but also the software beyond it that processes the messages received from thalamus and will then how the robotic head will move.

The SKENE component is only mentioned as part of the interaction towards the human players, nevertheless, it is used as a translator that receives Skene Utterances, which will be mentioned later what they are, from our agent architecture with the objective to process them and send events directly to EMYS, which will then produce the interactions relatively to those Skene Utterances received.

B. Agent's Architecture

Our Agent's Architecture is divided in different modules, where every single one of them as certain input that depends on the output generated by the previous module, and produces an output to provide for the next module.

1) *Perception Receiver*: In the perception receiver module, the grand objective is to receive the various perceptions that are meant for the agent, be those perceptions events from the game, which could be actions from the players, including the agent, and events to make the agent play, or events sent from the other components associated with all the system, for example, when a the robot as stopped animating or a speech as ended.

2) *Memory*: This component holds all the events and game states that the agent as perceived, which does not include information created by the agent, such as, it's knowledge about the other players.

3) *Theory of Mind*: Receive the information from the Memory component about what has happened and what

is the game state and calculates the probability of each player having a certain card and the probability of doing a certain action. This information is then given to the Decision Making Algorithm component. These probabilities mentioned before towards the agent itself, are not calculated through this component, but are received from the Decision Making Algorithm component and kept here so that the same type of information is kept on the same module.

4) Decision Making Algorithm: The Decision Making Algorithm is the module where the actions performed by the agent are generated, this also includes the interactions done, that require the robot, to interact with the other players in the real world.

5) Action/Interaction Producer: In this final component, the information received, as mentioned before, is the action produced by the Decision Making Algorithm component. This action is then processed alongside with what agent knows about the game to produce a possible interaction that will be sent to the robot in a way to give the agent a social presence and with that be capable of producing behaviour that is deceptive.

C. Digital Tabletop and Coup Unity Game

The solution for this, as mentioned before on the Problem/Solution chapter, is a digital tabletop that is big enough to allow multiple players and is able to run software.

The game had to be built in a way that allows the modification of its own state through the use of a virtual interface as well as a physical interface. For the virtual interface we used a mechanism that receives events from a central messaging system called Thalamus. These events would then be applied unto the game state, modifying it. For every action or modification of the game state, the Coup Unity game sends an event to the central messaging system, Thalamus, that, when applied to the previous game state, produces the current game state.

Regarding the interface that allows the humans to interact with the game, it is graphically represented through the digital tabletop. Each of the players, including our virtual player, will have a determined position where all their information is. This information includes their cards, that are hidden by default, a button that shows or hides the cards, depending on their current state, hidden or shown, a number that represents their current number of coins and finally, a list of actions, much like the original summary card, that not only allows the players to do the actions, but also provides all the information they need in terms of actions, counteractions and which characters are needed for any of those.

D. Robot EMYS

Our chosen robot to fulfill the requirements for the physical entity was EMYS (EMotive headY System), which is an emotive robotic head designed and built within the EU FP7 LIREC project. This head is composed of three discs and

equipped with a pair of eyes and eyelids that are movable. Everything is mounted on a movable neck[13].

The head is capable of speech through the use of a speaker, and is able to produce prerecorded or synthesized voices.

By using an already developed system to perform interactions through EMYS, we are able to easily and effectively make our agent interact with the human players, making this interaction as complex or simple as we want, do to its easiness of use.

The component that produces the interactions is called Skene, which is a semi-autonomous behaviour planner capable of semi-automated behaviour (cite SKENE). To produce such semi-automated behaviours, Skene takes as input a high-level behaviour description language that was developed by a team that also non-technical partners from psychology, which is called Skene Utterances, and perception information, such as target locations. The output of Skene consists on both the scheduling of BML (Behaviour Markup Language) and non-BML actions (such as sounds or application commands). This will then be sent to the EMYS component which in turn makes the actual robotic head move and reproduce sound according to the Skene Utterance.

VII. EXPERIMENTS

In order to test our virtual agent so that we can prove that our hypothesis is correct, we have designed a use-centered study with people playing the Coup game against our virtual agent.

The aim of the experiment, as mentioned before, is to see if our virtual coup agent is capable of playing at the same level of human beings, which include not only deceiving them, but also a small amount of discovering when the humans are lying.

The equipment we used to perform this experiment was:

- **A MultiTaction Ultra Thin Bezel Display**, which is a 55" display unit with interactive multiuser LCD, capable of tracking an unlimited amount of touch points, including hands, fingers, fingertips, 2D markers and real-life objects, with object recognition, as our digital tabletop;
- **An EMYS** (EMotive headY System), as our robot that physically represents our agent and interacts with the human players;
- **Three Lavalier microphones** to record the human player voices;
- **Four cameras** for filming, where one had the sole objective to record the interaction of the agent with the human players, and the other three focused on the behaviour of each of the human players;
- **One Coup game set** to explain the game to participants that did not know how to play it, or did not remember it so well.

A. Procedure

Regarding the sample, a total of 57 university students took part of this study, where 38 were male and 19 female, with ages ranging from 19 to 29.

Upon arrival, participants were allocated to just one of the types of group (playing individually against our agent or in a group with other two human players and the agent) and conditions (the lie condition, the no lie condition). Participants were not aware that the lie/no lie condition existed, so their initial perception of our agent did not differ between those two conditions.

They started by filling a pre questionnaire without supervision and then, after finishing it, the game Coup was explained using the original board game, so that all the players had at least a basic understanding of the game and were capable of playing it.

After the game session, participants were taken to a different room and then filled a pos questionnaire, without being supervised, in which the questions were in regard on how they felt about the interaction with our agent, more specifically, through the use of the EMYS robot as our agent physical representation.

The participants were then thanked for their participation in our experiment and contributing for our study, and were gifted with a coupon to get a free ticket for any movie in a certain group of cinemas.

After the interaction with our agent, by playing the game, the participants filled another questionnaire, this time being the Godspeed Questionnaire [2], in order to understand if the perception of the robot changed regarding the condition that they were allocated to. For this, a Credibility scale was also applied (taken only the Trustworthiness dimension from Ohanian, 1990) in order to understand if participants felt when our agent was being dishonest (all this was answered in a 5-point Likert scale) [14]. A Trust scale specific for Human-Robot Interaction was then used to perceive the level of trust the participants had on our agent, through its physical representation with EMYS [15].

In the last questionnaire, participants also answered directly to what they thought of our agent. They were asked on what level they would put the Coup playing skill of our agent, how much they thought that the agent lied, and how well it lied. All this questions were answered by using a 5-point scale.

Other measures that we used to understand how well our agent was capable of deceiving and ultimately winning the game against human players were the percentage of victories that the agent achieved, the number of times that he got challenged and won, and finally, the number of challenges he got per session. The first measure was used to know if in reality, an agent that has the capability of lying is indeed more beneficial than an agent that only tells the truth, the second measure to know if the agent was capable of successfully deceiving its human opponents, the last measure was needed to know if the participants were getting more doubtful of the agent or not.

B. Measures

To understand how trust would be perceived by a human player regarding a robot, by playing the deceptive game of Coup, two questionnaires were used.

Before the game session, participants responded to the Big Five Questionnaire [10] to ascertain the participant personality type (validated for the Portuguese population by Lima and Castro 2009), followed by an interpersonal trust scale, the Multidimensional Trust Scale [4] to see the level of trust that the participants had in themselves and others. For this scale, only the dimensions of Self and Others to ascertain the global score of trust were used, leaving the Environment dimension out due to its low internal consistency value.

C. Results

The results we got regarding the Multidimensional Trust Scale had the objective to prove that our participants, having been attributed to different group types and conditions, were still a good sample in terms of their initial trust value.

The participants from both group type conditions have very similar values for the trust value, not only that, but within both group types, even for the different lying conditions (lie and truth) the means are practically identical.

Continuing with the results obtained from the Godspeed Questionnaire, we found out that all of the measures of perceptions that we captured from the participants have an increase from the Group condition to the Individual condition, being the Likeability measure the one with the biggest increase for the Lie condition, an increase of perception from 3,08 to 3,56, or in other words, and increase of 10% in a 5-point Likert scale, which in itself represents that participants perceived our agent to be 10% more likeable in the Individual condition than in the Group condition. Other conclusion that we can reach based on these results is that, given that the game is played by multiple players and since the game also has in it player exclusion, by being removed earlier from the game, our agent's interaction, and consequently how it is perceived by the participants, diminishes as he will no longer take turns and perform actions.

Still in the results from the Godspeed Questionnaire, we can see differences between the Lie and Truth conditions. Focusing primarily in the Individual condition, as the differences are more significative, we can see that the participants in the Lie conditions perceived our agent as being more Anthropomorphic, Animate, Likeable, Intelligent and more dishonest based on the Credibility score, since the more score it has in the Credibility measure, the more dishonest its behaviour was. Taking this into account, we can come to the conclusions that by being capable of deceiving its fellow opponents, these perceived it as being more similar to the average human being and with that thought of it as more human, giving a better averaged score to every single measure in terms of this questionnaire.

In terms of the results obtained from the Trust scale specific for Human-Robot Interaction, we got a very similar percentage of truth in the Group condition for both Lie and Truth conditions, with around 60% trust from the participants towards our agent, this may be due to the fact that its interaction on this group type condition did not have the most impact on the participants, so they felt the same towards our agent int both the lying conditions. More interesting,

even though not big enough to be scientific significant, is the difference between the Lie and Truth conditions in the Individual group type condition, having an increase of 7,53% of trust from the Truth condition to the Lie condition. This increase is really interesting since the participants knew at one point or another that our agent was capable of deceiving, so why would they attribute an higher score of trust to a deceiving agent? The reason to that is most likely the same for the increased interaction perception from the participants in this condition, the Individual Lie condition, which is the fact that the participants more easily identify themselves with our agent and consequently build a more trustful image of it, perceiving it as more human.

Going into the results obtained from the more direct questions asked, the ones that only needed direct processing to build the results, we will start with how well the participants perceived our agent of playing. This measure was score with a 5-point Likert scale, so the opinion of each participant may differ slightly even if they give the same score as another participant, nevertheless, the results obtained here are a little different from the results obtained in other of various variables we used, in the sense that the conclusions that we can take from them are completely different based on the group type condition. In the Individual condition, we can see an 8% increase on the perception towards how well our agent played when comparing from the Truth condition to the Lie condition, going from an average score of 3,67 to 4,07. On the other hand, in the Group condition, we can clearly see that the average score given has decreased from the Truth condition to the Lie condition, going from an average score of 4,00 to 3,80, which represents a 4% decrease. This small decrease comes most likely from the fact that when playing in a group, sometimes a good strategy is one where the player will not try to get ahead in the beginning of the game, or in other words, will not be a potential threat in the beginning, making itself go unnoticed and avoid being the target of the other players, and since our agent in the Truth condition tends to play in a more conservative way, as its actions are limited, its quite plausible for that to be the reason of such decrease in the perception of how well it played when going from the Truth to the Lie condition.

Going a little deeper towards how well our agent played in the different conditions, we can see the results we got from the in game victories itself. We got that in terms of the Individual group type condition, in the Truth condition our agent won almost half of the games, with a 49,00% win ratio. This percentage receives a considerate increase that, while not being scientifically significant, contributes enormously to our perception of how well the agent played. The increase is of 7,00% to reach a much better win ratio of 56%, which in other words means that our agent is actually more capable of winning at the game of Coup than a human player. In terms of the Group condition, the results are a little exchanged for both the Lie and Truth condition, as have been shown in the last measures, being it an average of 23,40% win ratio for the Lie condition and an average of 25,00% win ratio for the Truth condition.

The conclusions we can take from the win ratio for the different conditions is that in all of them, our agent is capable of playing at least on the same level of a human player. In the Lie condition for the Individual group type condition, our agent can even play slightly better than a human, while in all other conditions, specifically in the Truth condition for the same group type condition, it won practically half of the games it played, where there was only one opponent, so the human players got the other half of games won, an in both the Lie and Truth conditions for the Group group type condition, the percentage of games won is close, on average, to one quarter of games won, which is actually the same average win ratio for each of the human players, since there are more three player when not counting with our agent. These are quite good results that we achieved, for both our lying and truthful agent, since they both use our Decision Making algorithm, we can prove that it is indeed possible to make an agent, that by implementing it on a social robot, is capable of playing at the same level of humans.

To see how having both the Truth and Lie conditions would influence the amount of times that our agent got challenged, we counted all of the times where a challenge was presented to our agent. The amount of times that our agent gets challenged increases significantly from the Truth condition to the Lie condition, with a 4 times increase in the Group condition. This is due to the fact that in the Group condition is easier to make actions that may seem deceiving and with that, participants will label our agent as being capable of deception and then start to challenge it. Regarding the Individual condition, the increase while not being as significative as the one in the Group condition, is still a 50% increase. This leads us to conclude that the participants in general, knew that our agent in the Lie condition was more deceitful than our agent in the Truth condition, and with that, they challenged it more.

Finally, the last results that we analyzed were the average number of times per session that our agent when challenged, had the card to actually produce the action that he got challenged on. The results show exactly that, with our agent having relatively the same amount of times that it won in the Individual condition when challenged, but having a significant increase from the Truth to the Lie condition in the Group condition. As mentioned before, having more people that detect that it is capable of deceiving, will make the amount of challenges he receives increase, and by capitalizing on it, by starting to play a little more truthfully, our agent is capable of winning a great amount of challenges that it receives. This on the other side can have some repercussions, as the other players start to lose their cards, they will start to consider our agent as a potential threat, since it was our agent that made them lose their cards and consequently being close to losing, justifying why the percentage of won games for our agent in the Group Lie condition is inferior to the one of the Group Truth condition.

VIII. CONCLUSIONS

In this document we started by presenting some background on Deception, Board Games and the COUP game, with an extensive explanation on how the game works, in order to provide some comprehension towards some topics that are not usually directly linked to the computer science area. We then presented some of the most important works that relate to our own work's state of the art and consequently contribute to the ideas developed here. These works range from agents that can act deceptively in a purely digital environment, to agents that were implemented on robotic entities and were capable of deception. Other works that were focused simply tried to prove that a robot can actually deceive a human being and not just another robot.

We proceeded to define both the problem that surrounds our hypothesis and the way to solve it, just focusing on the components needed to do so in a more general and abstract way. After having the problem and solution defined, we presented the decision making algorithm that would be the base of the action selection done by our agent, before that, we introduced the algorithm that was the base for our algorithm and why minimizing counterfactual regret would help us solve our problem and successfully prove our hypothesis. The presentation of our agent architecture plus all other components needed to make the experiences to prove our hypothesis was done in a chapter specifically for just the implementation of our solution. In this chapter, the general solution that we presented before was specified into a real solution that we could effectively use in the real world.

The experiment done was then presented and included all the information regarding the sample we used, all its procedure, the variables that we used to measure different things, the results we obtained from analyzing the data acquired during it, a discussion towards the results we got, taking some conclusions from most of the results and finally concluding with what was the most interesting result we got from the experiment and what could we have done better to obtain more significant results towards the interaction between our agent and the participants.

The result that was expected to give the most feedback to the way our work was going was definitely the percentage of games won by our agent against the human players. As already mentioned before, our agent was capable of achieving a 50% win ratio against a single human opponent and a 25% win ratio against three other human players, which in itself means that our agent was playing at the same level of humans. This completely proves our hypothesis, which is the greatest achievement that this work could get, a successfully proven hypothesis. Not only could our agent play at the same level of humans, but by using the unrestricted decision making algorithm, the one that is capable of using actions that were lies, our agents was capable of playing at a slightly higher level than its human opponents.

REFERENCES

- [1] Simon Baron-Cohen. *Mindblindness: An essay on autism and theory of mind*. MIT press, 1997.
- [2] Christoph Bartneck, Dana Kulic, Elizabeth Croft, and Susana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 1(1):71–81, 2009.
- [3] Jr. Bond, CharlesF. and Michael Robinson. The evolution of deception. *Journal of Nonverbal Behavior*, 12(4):295–307, 1988.
- [4] Karen Carrington. Toward the development of a new multidimensional trust scale. 2007.
- [5] Cristiano Castelfranchi, Rino Falcone, and Fiorella De Rosis. Deceiving in golem: how to strategically pilfer help. Citeseer.
- [6] Dorothy L Cheney and Robert M Seyfarth. *Baboon metaphysics: the evolution of a social mind*. University of Chicago Press, 2008.
- [7] David Christian and R Michael Young. Strategic deception in agents. Master's thesis, Graduate Faculty of North Carolina State University, 2004.
- [8] Bella M DePaulo, Deborah A Kashy, Susan E Kirkendol, Melissa M Wyer, and Jennifer A Epstein. Lying in everyday life. *Journal of personality and social psychology*, 70(5):979, 1996.
- [9] Joao Dias, Ruth Aylett, Henrique Reis, and Ana Paiva. The great deceivers: Virtual agents and believable lies.
- [10] Samuel D Gosling, Peter J Rentfrow, and William B Swann. A very brief measure of the big-five personality domains. *Journal of Research in personality*, 37(6):504–528, 2003.
- [11] Maaike Harbers, Karel van den Bosch, and John-Jules Meyer. Modeling agents with a theory of mind. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 217–224. IEEE Computer Society, 2009.
- [12] M. D. Hauser. Costs of Deception: Cheaters are Punished in Rhesus Monkeys (*Macaca mulatta*). *Proceedings of The National Academy of Sciences*, 89:12137–12139, 1992.
- [13] Jan Kedzierski, Robert Muszyński, Carsten Zoll, Adam Oleksy, and Mirela Frontkiewicz. Emysemotive head of a social robot. *International Journal of Social Robotics*, 5(2):237–249, 2013.
- [14] Roobina Ohanian. Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of advertising*, pages 39–52, 1990.
- [15] Kristin E Schaefer. *The perception and measurement of human-robot trust*. PhD thesis, University of Central Florida Orlando, Florida, 2013.
- [16] Kazunori Terada and Akira Ito. Can a robot deceive humans? In *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on*, pages 191–192. IEEE, 2010.
- [17] Eugenio Di Tullio Patricia Alves-Oliveira Ana Paiva Tiago Ribeiro, André Pereira. From thalamus to skene: High-level behaviour planning and managing for mixed-reality characters. 2014.
- [18] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. In *Advances in neural information processing systems*, pages 1729–1736, 2007.